FIG. 1

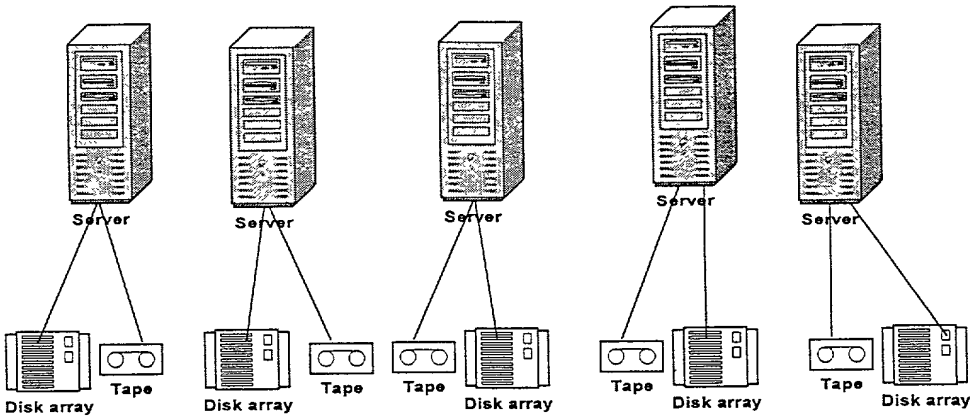
Server
Server
Server
Server
Server
Tape
Tape
Tape
Tape
Tape
Disk array
Disk array
Disk array
Disk array
Disk array

FIG. 2

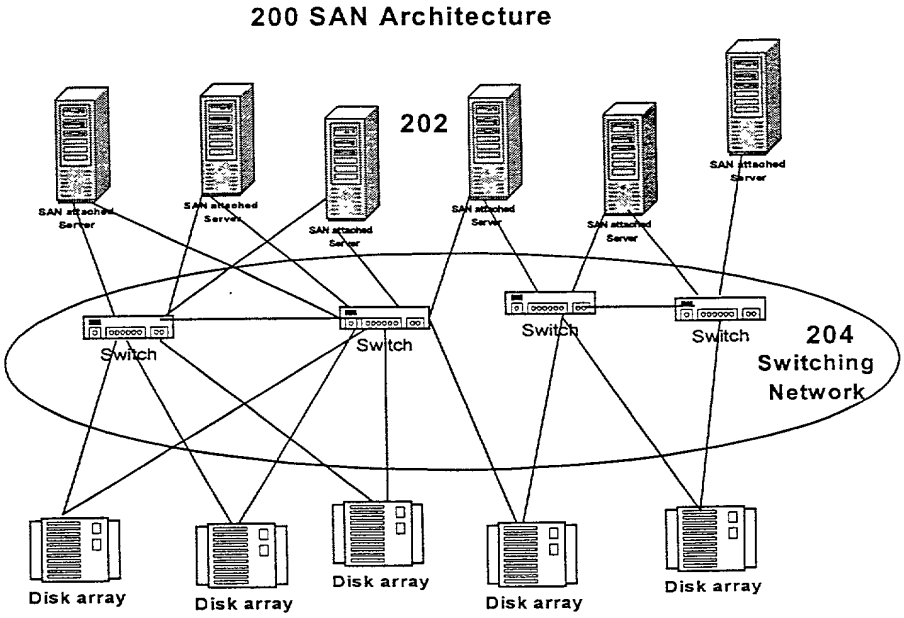
200 SAN Architecture
202
SAN attached Server
SAN attached Server
SAN attached Server
SAN attached Server
SAN attached Server
SAN attached Server
Switch
Switch
Switch
Switch
204
Switching Network
Disk array
Disk array
Disk array
Disk array
Disk array
206

FIG. 3

Steps to manually provision storage for a server (Prior Art):

STEP 251
A SAN administrator physically verifies that the server is connected to the storage subsystem by checking the fibre channel cabling from the server host bus adapter (HBA) to the fibre channel switch(es) and on to the target disk subsystem.

STEP 252
A SAN administrator uses a storage subsystem tool (or SAN management tool) to find a data volume of sufficient size.

STEP 253
A SAN administrator uses a storage subsystem tool to set mapping of the selected data volume to the desired host.

STEP 254
A SAN administrator makes a best guess as to the usage of various switched fabrics and paths and selects one or more paths from the server to the data volume.

STEP 255
A SAN administrator uses the SAN switch tool to verify that the physical connection is good by checking the name server data to see that the server (HBA) has successfully logged into the switch.

STEP 256
A SAN administrator uses the SAN switch tool to configure the fabric by setting the switch zones between the storage port and server HBA port.

STEP 257
Step 255 and 256 are repeated by the SAN administrator for every fabric used between the server and the data volume.

STEP 258
The SAN administrator uses the HBA interface tool to configure the HBA's on the server to set masking and mapping for the data volume that the server needs to access.

FIG. 4

Block Diagram illustrating a system for performing Automated Provisioning within a SAN

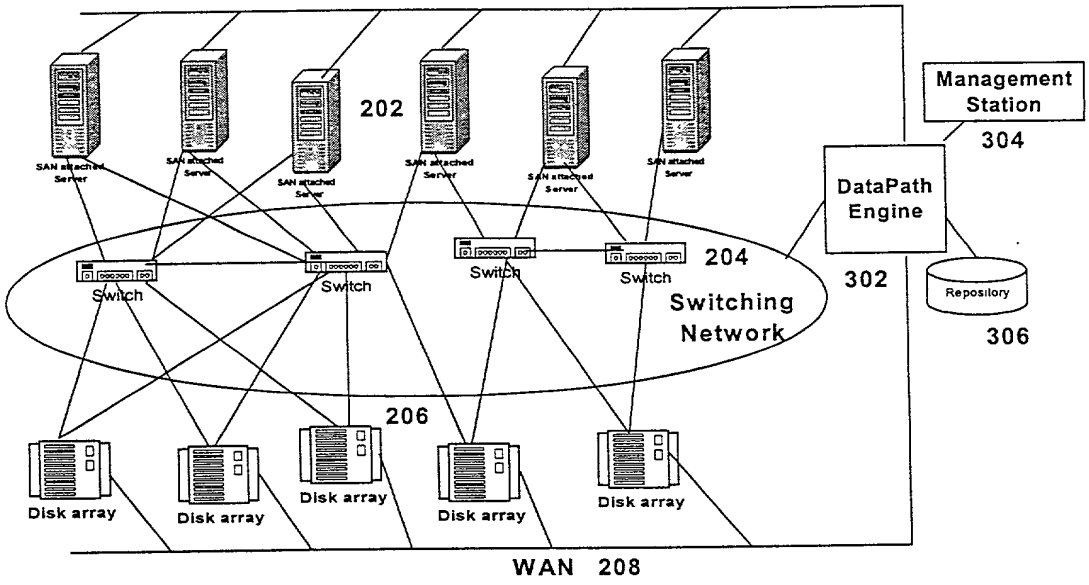

FIG. 5

Steps for Automated Provisioning of Data Paths:

STEP 351
DataPath Engine software logic discovers all physical devices including switches, storage subsystems, servers, HBAs, data volumes and SAN appliances; all connections through the SAN including switch connections, server HBA connections and storage subsystem connections; all settings including switch zones, data volume maps, security settings, and SAN state; and displays the graphical topology of the SAN.

STEP 352
An Operator selects a create data path command from a graphical interface. The graphical interface provides a wizard that guides the operator through the data path creation and requests the operator to select an application and size of data volume (or existing data volume). Selection rules and policy is pre-created for each application.

STEP 353
DataPath Engine 302 selects the best data volume that meets the policy criteria. Best selection pathing methodology takes into consideration the state and activity of the SAN. DataPath Engine 302 discovers and selects physical paths through the SAN fabric from the selected data volume to the target application host that meet the criteria. Best selection pathing methodologies select the best path based upon the state and usage of the SAN.

STEP 354
The selected data paths are displayed for the operator. If the operator accepts the selection, DataPath Engine 302 configures settings on the HBA, storage device, data volume, and fabric switches to maximize security and allows accessibility to the data volume. The path settings and configuration data are saved.

FIG. 6

Example Rules that can be used for Data Path Selection criteria

EXAMPLE RULES	DESCRIPTION	EXAMPLE VALUES
451: Storage Type	The end point storage type	Disk, Tape, SSD, etc.
452: Number of Threads or Channels	Sets the data path rule to the number different physical connections between server and data Volume that must be selected. Each Thread requires a different HBA port on the Application Server.	1-10
453: Number of Fabrics or WANs	Sets the rule to the number of different switch fabrics to use for selection of each Thread above. Each fabric requires a different HBA port in the Application Server.	1-10
454: Bandwidth	Sets the rule as to whether each Thread selected must be exclusive or shared. Exclusive Threads mean that no other Thread may share any portion of the physical connection and ports used.	Shared or Exclusive
455: Data Volume Security	Sets the rule as to whether to select a storage subsystem where Data Volume mapping is used or not. Open does not use storage subsystems with Volume Mapping.	Secure or Open
456: Data Volume Action	Sets the policy for action if a Data Volume map is modified by alternate SAN tools: Ignore (default for Open), Notify only, Notify and disable by hiding the data Volume from server, or Notify and override (change it back to original settings)	Ignore, Notify, Disable, Override
457: Fabric Security	Sets the rule as to whether to select a switch with zoning enabled or not. Open does not use switches with zones.	Secure, Notify or Open
458: Fabric Action	Sets the policy for action if a switch zone is modified by alternate SAN tools: Ignore (default for Open), Notify only, Notify and disable by hiding Volume from server, or Notify and override (change it back to original settings)	Ignore, Notify, Disable, Override
459: HBA Security	Sets the rule as to whether HBA masking or mapping is used or not. Open does not use HBA security settings	Open, Secure
460: HBA Action	Sets the policy for action if a HBA setting is modified by non-InterSAN software. (same as Fabric Action)	Ignore, Notify, Disable, Override
461: Data Volume Size	Number of Gigabytes. The closest Volume size will be discovered.	1-99999
462: Volume Pooling	Use more than one volume to meet size requirements	Yes/no
463: Data Volume Char	Type of RAID characteristics of the data volume	RAID 0,1,3,5, etc
464: Data Path Stat	Status of data path: Enabled or Disabled	Enable or Disable

Fig. 7

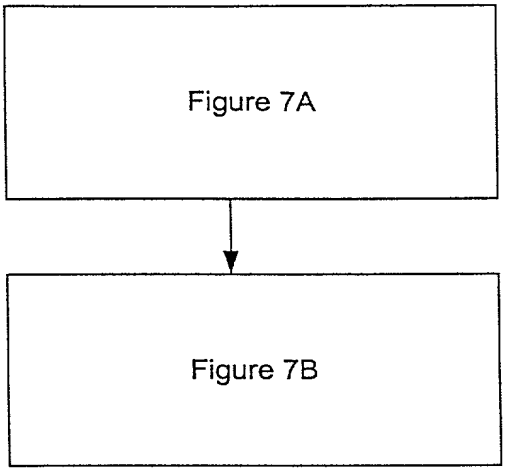

FIG. 7A

Pathing Methodology for optimal data path selection (Process 550)

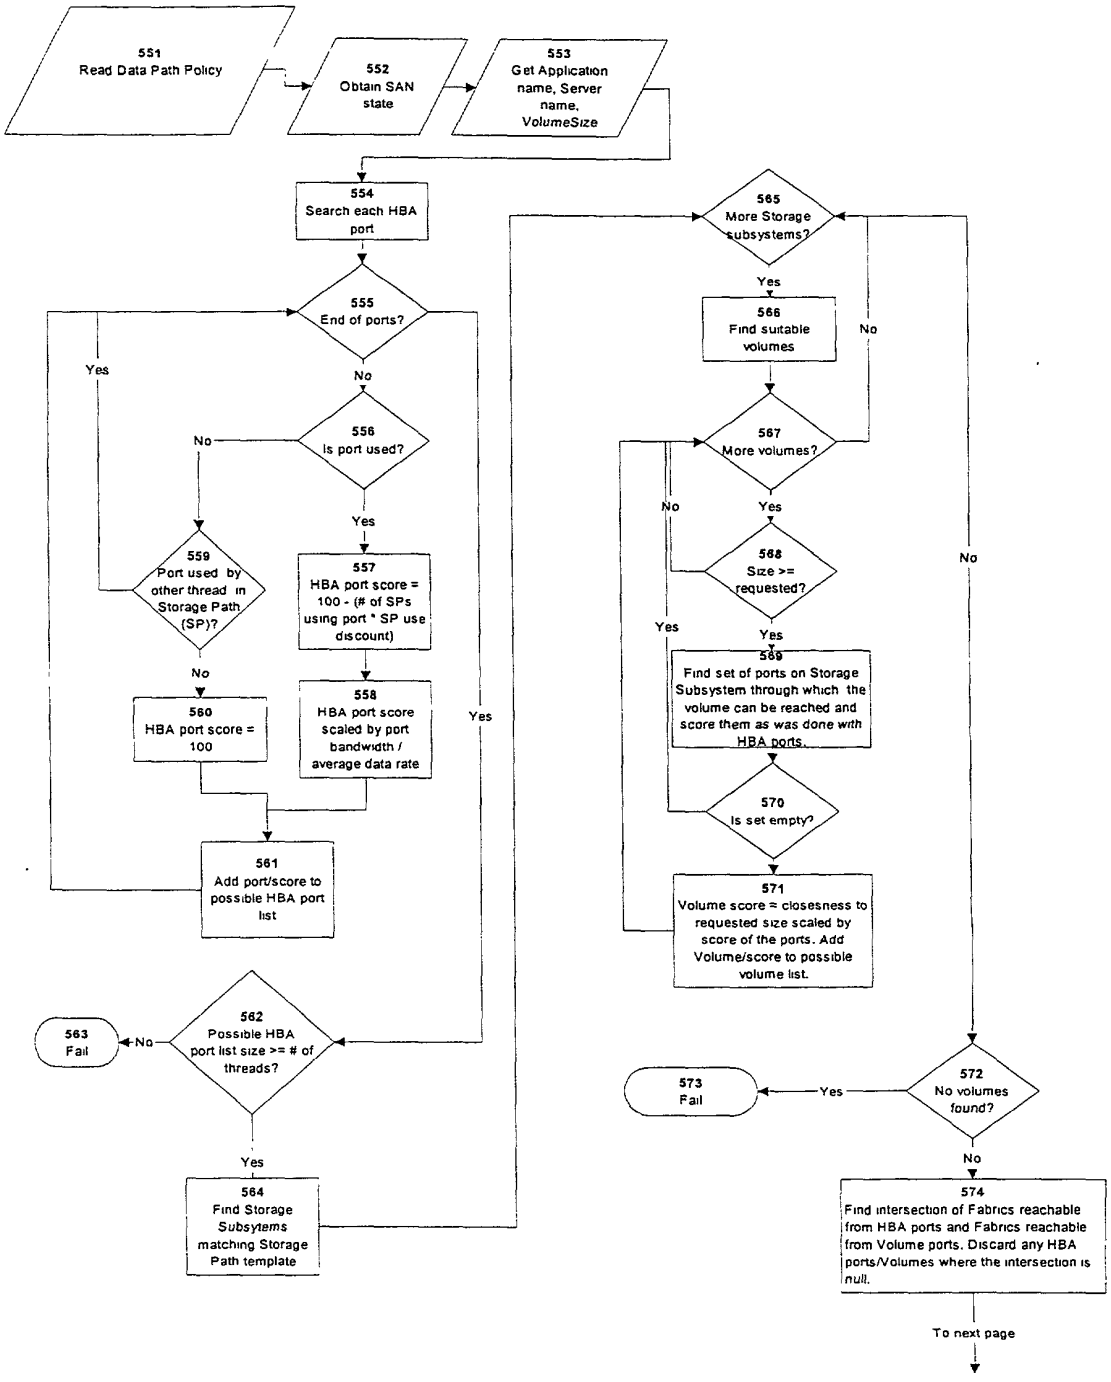

FIG. 7B

Pathing Methodology for optimal data path selection (continuted)

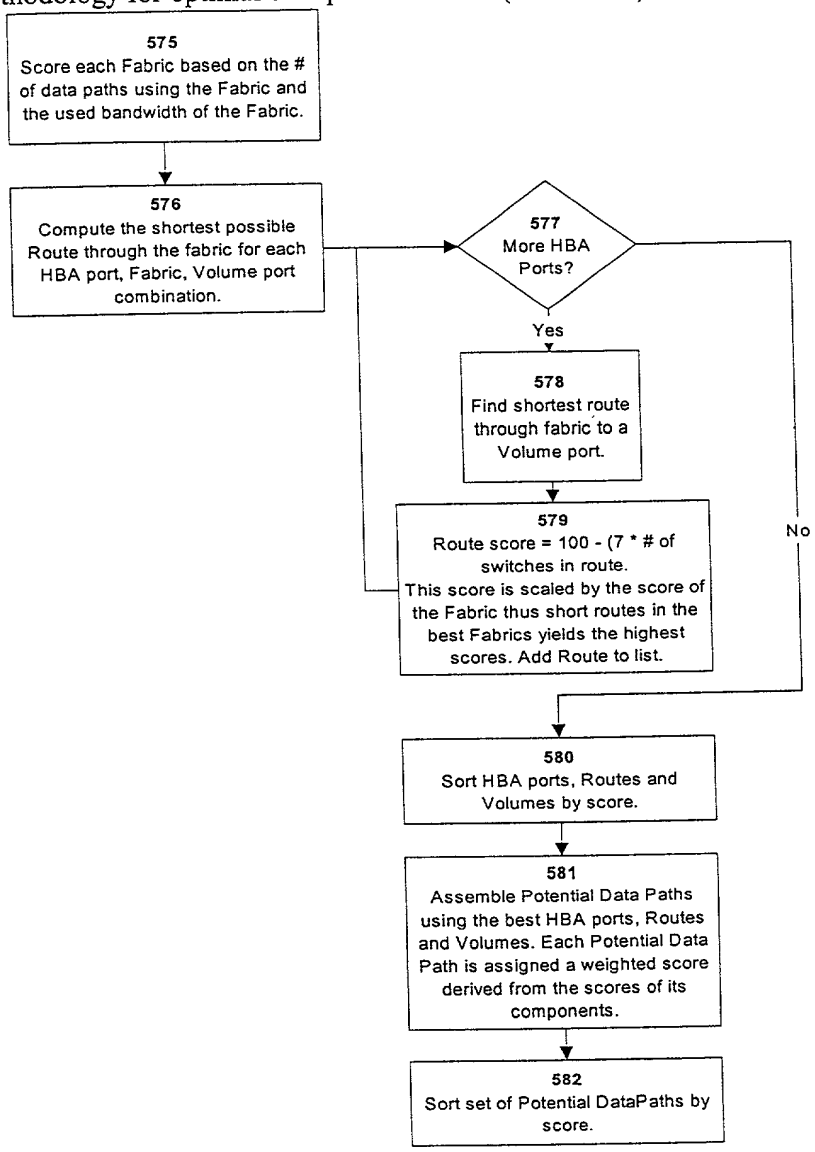

FIG. 8

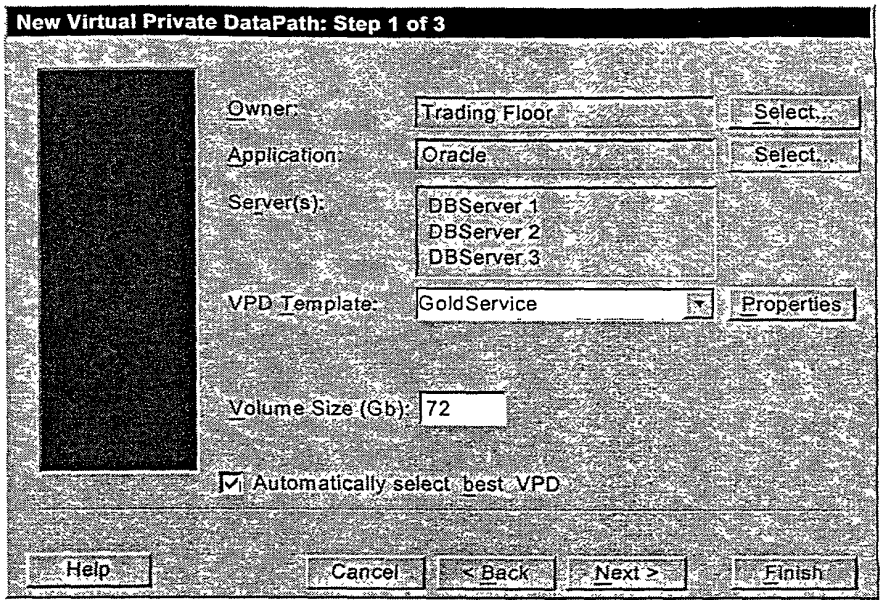
New Virtual Private DataPath: Step 1 of 3
Owner:
Trading Floor
Select...
Application:
Oracle
Select...
Server(s):
DBServer 1
DBServer 2
DBServer 3
VPD Template:
GoldService
Properties
Volume Size (Gb):
72
Automatically select best VPD
Help
Cancel
< Back
Next >
Finish

FIG. 9

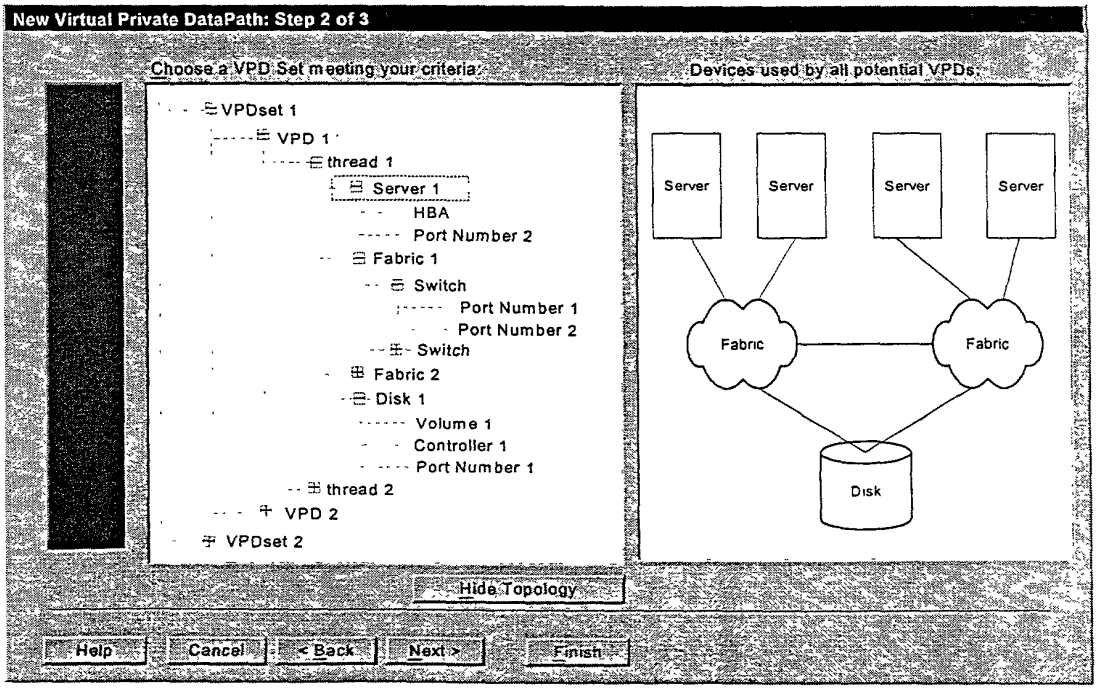
New Virtual Private DataPath: Step 2 of 3
Choose a VPD Set meeting your criteria:
VPDset 1
VPD 1
thread 1
Server 1
HBA
Port Number 2
Fabric 1
Switch
Port Number 1
Port Number 2
Switch
Fabric 2
Disk 1
Volume 1
Controller 1
Port Number 1
thread 2
VPD 2
VPDset 2
Devices used by all potential VPDs:
Server
Server
Server
Server
Fabric
Fabric
Disk
Hide Topology
Help
Cancel
< Back
Next >
Finish

FIG. 10

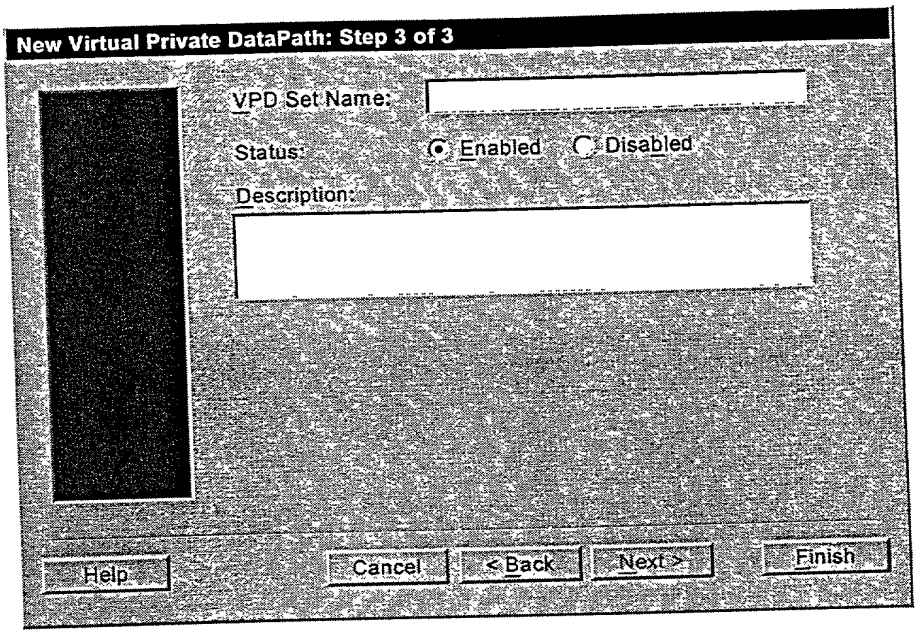
New Virtual Private DataPath: Step 3 of 3
VPD Set Name:
Status:
Enabled
Disabled
Description:
Help
Cancel
< Back
Next >
Finish

FIG. 11

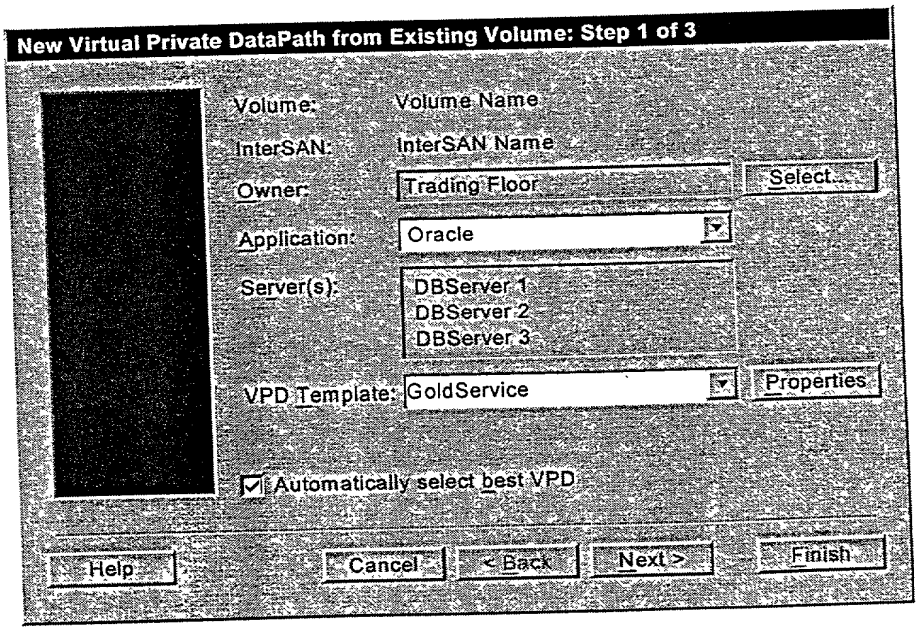
New Virtual Private DataPath from Existing Volume: Step 1 of 3
Volume:
Volume Name
InterSAN:
InterSAN Name
Owner:
Trading Floor
Select...
Application:
Oracle
Server(s):
DBServer 1
DBServer 2
DBServer 3
VPD Template:
GoldService
Properties
Automatically select best VPD
Help
Cancel
< Back
Next >
Finish